

(12) UK Patent Application (19) GB (11) 2 308 774 (13) A

(43) Date of A Publication 02.07.1997

(21) Application No 9706681.5

(22) Date of Filing 11.06.1993

Date Lodged 02.04.1997

(30) Priority Data

(31) 9307442

(32) 08.04.1993

(33) GB

(62) Derived from Application No. 9312128.3 under Section 15(4) of the Patents Act 1977

(51) INT CL⁶
H04N 5/14

(52) UK CL (Edition O)
H4F FD10 FD12K9 FD12X FD30J FD30K FD30R FD54
FGM

(56) Documents Cited
GB 2231749 A

(58) Field of Search
UK CL (Edition O) H4F FEP FER FEX FGM
INT CL⁶ G06T 7/20 , H04N 5/14 7/01
Online database: WPI

(71) Applicant(s)

Sony United Kingdom Limited

(Incorporated in the United Kingdom)

The Heights, Brooklands, WEYBRIDGE, Surrey,
KT13 0XW, United Kingdom

(74) Agent and/or Address for Service

D Young & Co
21 New Fetter Lane, LONDON, EC4A 1DA,
United Kingdom

(72) Inventor(s)


Morgan William Amos David
Martin Rex Dorricott
Shima Ravji Varsani

(54) Selection of global motion vectors for video signal processing


(57) Motion compensated video signal processing apparatus includes means (Fig 1) for generating local motion vectors representative of the motion in respective blocks of two images of a pair, means 430 for deriving from these local motion vectors a set of global motion vectors comprising the most commonly occurring, distinct local motion vectors and means 440 for selecting those global motion vectors which occur more frequently than a global threshold value. A switch 450 outputs for each block a group of motion vectors comprising the zero motion vector, the local motion vector for that block and the global motion vectors.

A frequency array 420 accumulates the number of occurrences of each local motion vector and an array scanner 430 selects those which differ from the next higher selected vector by more than the difference threshold 435. An operator may vary the global threshold which is set, for example, at 0.5% of the total number of local motion vectors derived from the image pair. Preferably 8 unique global motion vectors are output for each block.

FIG. 3



GB 2 308 774 A







FIG. 3









FIG. 7



MOTION COMPENSATED VIDEO SIGNAL PROCESSING

This invention relates to motion compensated video signal processing.

This is a divisional application from the parent application 9312128.3.

5 Motion compensated video signal processing is used in applications such as television standards conversion, film standards conversion and conversion between video and film standards.

10 In a motion compensated television standards converter, such as the converter described in UK patent application number GB-A-2 231 749, pairs of successive input images are processed to generate sets of motion vectors representing image motion between the pair of input images. The processing is carried out on discrete blocks of the images, so that each motion vector represents the inter-image motion of the content of a respective block.

15 In the motion vector estimation process a correlation surface representing spatial correlation between blocks of the two input images is examined to detect points of maximum correlation. (The correlation surface actually represents the difference between the two input images so that the points of maximum correlation are in fact minima on the correlation surface, and will be referred to as such herein.) When a minimum has been detected, a motion vector is generated from the spatial
20 position of the minimum in the correlation surface. Tests are performed to determine whether that minimum represents a significant peak in correlation with respect to the remainder of the correlation surface. If the minimum passes this test it is treated as being "valid", and a confidence flag associated with the motion vector is set.

25 After the individual motion vectors, termed "local" motion vectors, have been derived for the blocks of the input image, the local motion vectors are supplied to a motion vector reducer. The motion vector reducer assigns to each block a set of motion vectors selected, in order of preference, from: the zero motion vector; the local motion vector for that block; the local motion vectors for blocks neighbouring that block in the input image; and a set of vectors known as "global" motion vectors.
30 The global motion vectors are derived for each image by ranking all the (valid) local motion vectors for this image in order of decreasing frequency, and then selecting from the most common of the local motion vectors a number of unique motion

vectors to be used as global motion vectors. A constraint is applied here to ensure that the global motion vectors differ from each other by at least a minimum amount.

5 The sets of motion vectors for all the blocks of the input image are then supplied to a motion vector selector. The purpose of the motion vector selector is to allocate to each picture element (pixel) of a block of an output image a single motion vector selected from the set of motion vectors supplied for a corresponding block of the input image. As this motion vector will be used in interpolation of the output pixel, the set of motion vectors supplied to the motion vector selector are tested, by performing a correlation test on blocks of pixels pointed to by each of the set of motion vectors, so that the "best" motion vector is selected for each pixel from the set of motion vectors for the corresponding block. The selected motion vector for each pixel is then supplied to a motion compensated interpolator which uses these motion vectors to interpolate the pixels of an output image from the pair of input images in dependence upon the temporal offset of the output image from the pair of input images.

15 The purpose of the motion vector reducer is therefore to provide the motion vector selector with a reasonable selection of motion vectors for each block from which to select the most appropriate motion vector for each output pixel corresponding to that block. The global motion vectors each represent a motion which was frequently detected in the image as a whole, so that it is often reasonable to include these vectors in the sets of vectors from which the most appropriate vectors are selected by the motion vector selector.

20 It has been observed that rapidly moving portions of a video image have a lower spatial detail than stationary or slowly moving portions. This means that even if a motion vector can be estimated which accurately represents the rapid motion of the object, the correlation test performed during vector selection (which relies on detecting correlation between parts of the same object in two successive input fields) can be unreliable.

25 This invention provides a motion compensated video signal processing apparatus comprising:

30 means for generating a plurality of motion vectors to represent image motion between a pair of input images of an input video signal, for use in interpolation of an

output image of an output video signal;

means for detecting whether each motion vector is greater than a predetermined size; and

5 means for restricting the use, in interpolation of the output image, of a motion vector detected to be greater than the predetermined size.

10 Accordingly, even though large magnitude (long) motion vectors (i.e. vectors representing rapid image motion) can be generated, these large vectors are deliberately restricted in their use in interpolation of the output image. This is because the lack of spatial detail in fast-moving objects (e.g. due to camera integration blur) means that:

a) the vector selection process applied to rapidly moving objects can be unreliable, and

15 b) for rapidly moving objects, the eye is less able to detect the judder introduced by non motion compensated interpolation; this judder can therefore be less subjectively disturbing than motion compensated interpolation using erroneously selected vectors.

20 In one preferred embodiment each motion vector comprises a plurality of coordinate values representing image motion along respective coordinate axes; and the means for detecting comprises means for comparing each of the coordinate values with a respective predetermined threshold value associated with each of the coordinate axes.

25 In another preferred embodiment each motion vector again comprises a plurality of coordinate values representing image motion along respective coordinate axes; and the means for detecting comprises means for generating a magnitude value from the coordinate values for each motion vector, the magnitude value representing the magnitude of the image motion represented by that motion vector, and means for comparing the magnitude value with a predetermined threshold value. In this case, it is preferred that the means for generating a magnitude value is operable to detect the square root of a weighted sum of the squares of the coordinate values. A
30 weighted sum may be used to reflect a different scaling of motion vector coordinate values in the various coordinate directions. Alternatively, the weighting could be applied equally to the individual squared coordinate values.

Although the coordinate axes could lie along various directions with respect to the input images, it is preferred that the coordinate axes represent vertical and horizontal image motion.

5 Preferably each motion vector represents image motion of a respective block of one input image of the pair between that image and the other image of the pair, and the apparatus comprises: means for deriving a set of global motion vectors comprising a plurality of distinct motion vectors selected from the most common of the plurality of motion vectors; and means for assigning to each of the blocks a group of motion vectors selected from the zero motion vector, the motion vector
10 representing image motion of that block and the global motion vectors.

In order to prevent global vectors being selected which only represent the motion of a small part of the image, it is preferred that each motion vector selected as a global vector has at least a predetermined number of occurrences among the plurality of motion vectors.

15 In one embodiment the "long" motion vectors can be prevented from being used as a global motion vector. To this end, it is preferred that the means for restricting comprises means for preventing selection as a global motion vector of a motion vector detected to be greater than the predetermined size.

20 In another embodiment, the use of the "long" motion vector can be further restricted, by employing a means for restricting comprising means, responsive to a detection that the motion vector representing motion of a block is greater than the predetermined size, for controlling the means for assigning to assign to that block only the zero motion vector and the motion vector representing motion of that block.

25 In a further preferred embodiment, the "long" motion vectors can be completely discarded, by employing a means for restricting comprising means for preventing the use, in interpolation of an output image, of a motion vector detected to be greater than the predetermined size.

30 In order that different measures can be taken for motion vectors in two different size ranges, it is preferred that the apparatus comprises: means for detecting whether each motion vector is greater than a second predetermined size, the second predetermined size being greater than the firstmentioned predetermined size; and means for preventing the use, in interpolation of the output image, of a motion vector

detected to be greater than the second predetermined size.

5 In another preferred embodiment the means for restricting comprises means for setting a flag, associated with each motion vector, to indicate whether that motion vector is greater than the predetermined size, and the apparatus comprises: means for selecting a motion vector, from the plurality of motion vectors, for use in interpolation of each pixel of the output image; a motion compensated pixel interpolator operable according to at least two modes of operation; and means for selecting one of the modes of operation of the motion compensated pixel interpolator, for interpolation of a pixel of the output image, in dependence on whether the flag, corresponding to the motion vector selected for use in interpolation of that pixel, is set. In this way, the interpolation operation can be matched to the motion characteristics of the video signal.

10 Although various arrangements of the pixel interpolator can be employed to give different interpolation operations in dependence on whether the flag is set, it is preferred that the pixel interpolator is operable, according to one of the modes of operation, to interpolate between blocks of each of the two images, pointed to by the selected motion vector, according to a combining ratio dependent upon the temporal position of the output image with respect to the pair of input images; and the pixel interpolator is operable, according to another mode of operation, to interpolate between blocks of each of the two images, pointed to by the zero motion vector, according to an equal combining ratio.

15 As mentioned above, a problem can arise when the global vectors are derived from a relatively small number of occurrences. For example, if the majority of the image is stationary, except for a small moving object, the majority of the motion vectors for that image would be close to zero. The vectors representing the stationary parts of the image may differ slightly due to noise or aliasing, but would still be too close to one another to be selected individually as global vectors. The next most commonly occurring vector, which would therefore be selected as a global vector, would be the vector derived from the small moving object. If this vector is passed to vector selection as one of the four vectors associated with a non-moving block of the image, it may be spuriously selected during vector selection. The use of a grossly incorrect vector in this way can result in visible artifacts being generated.

Viewed from a second aspect this invention provides a motion compensated video signal processing apparatus in which motion vectors are generated to represent image motion between a pair of input images of an input video signal, the apparatus comprising:

5 means for generating a plurality of local motion vectors to represent image motion of respective blocks of one input image of the pair between that image and the other image of the pair;

means for deriving from the local motion vectors a set of global motion vectors comprising a plurality of distinct motion vectors selected from the most
10 common of the local motion vectors, each global vector having at least a predetermined number of occurrences among the local motion vectors; and

means for assigning to each of the blocks a group of motion vectors selected from the zero motion vector, the local motion vector for that block and the global motion vectors.

15 This aspect of the invention addresses the problem of using global motion vectors derived from a relatively small number of occurrences, by preventing the selection as a global motion vector of a local vector which occurs fewer than a predetermined number of times.

Preferably the means for deriving comprises means for preventing the
20 selection, as global motion vectors, of two or more local motion vectors which differ from one another by less than a predetermined amount.

In a preferred embodiment, the predetermined number of occurrences can be adjusted under the control of an operator.

In an advantageously simple embodiment the apparatus comprises means for
25 storing a frequency array comprising a plurality of addressable memory locations, each possible value of the motion vectors representing the address of a respective memory location; means for incrementing a frequency value stored at the memory location corresponding to each of the plurality of motion vectors; and means for detecting those memory locations having the highest stored frequency values.

30 Apparatus according to the invention is particularly usefully employed in a television standards conversion apparatus.

Viewed from a third aspect this invention provides a method of motion

compensated video signal processing, the method comprising the steps of:

generating a plurality of motion vectors to represent image motion between a pair of input images of an input video signal, for use in interpolation of an output image of an output video signal;

5 detecting whether each motion vector is greater than a predetermined size; and restricting the use, in interpolation of the output image, of a motion vector detected to be greater than the predetermined size.

10 Viewed from a fourth aspect this invention provides a method of motion compensated video signal processing in which motion vectors are generated to represent image motion between a pair of input images of an input video signal, the method comprising the steps of:

generating a plurality of local motion vectors to represent image motion of respective blocks of one input image of the pair between that image and the other image of the pair;

15 deriving from the local motion vectors a set of global motion vectors comprising a plurality of distinct motion vectors selected from the most common of the local motion vectors, each global vector having at least a predetermined number of occurrences among the local motion vectors; and

20 assigning to each of the blocks a group of motion vectors selected from the zero motion vector, the local motion vector for that block and the global motion vectors.

The invention will now be described by way of example with reference to the accompanying drawings, throughout which like parts are referred to by like references, and in which:

25 Figure 1 is a schematic block diagram of a motion compensated television standards conversion apparatus;

Figure 2 is a schematic block diagram of a motion vector reducer;

Figure 3 is a schematic block diagram of a global vector detector; and

30 Figures 4 to 7 show modifications to the apparatus of Figure 1 to control the use of motion vectors in response to a long vector flag.

Figure 1 is a schematic block diagram of a motion compensated television standards conversion apparatus. The apparatus receives an input interlaced digital

video signal 50 (e.g. an 1125/60 2:1 high definition video signal (HDVS)) and generates an output interlaced digital video signal 60 (e.g. a 1250/50 2:1 signal).

5 The input video signal 50 is first supplied to an input buffer/packer 110. In the case of a conventional definition input signal, the input buffer/packer 110 formats the image data into a high definition (16:9 aspect ratio) format, padding with black pixels where necessary. For a HDVS input the input buffer/packer 110 merely provides buffering of the data.

10 The data are passed from the input buffer/packer 110 to a matrix circuit 120 in which (if necessary) the input video signal's colorimetry is converted to the standard "CCIR recommendation 601" (Y,Cr,Cb) colorimetry.

15 From the matrix circuit 120 the input video signal is passed to a time base changer and delay 130, and via a sub-sampler 170 to a subsampled time base changer and delay 180. The time base changer and delay 130 determines the temporal position of each field of the output video signal, and selects the two fields of the input video signal which are temporally closest to that output field for use in interpolating that output field. For each field of the output video signal, the two input fields selected by the time base changer are appropriately delayed before being supplied to an interpolator 140 in which that output field is interpolated. A control signal t , indicating the temporal position of each output field with respect to the two selected input fields, is supplied from the time base changer and delay 130 to the interpolator 140.

20 The subsampled time base changer and delay 180 operates in a similar manner, but using spatially subsampled video supplied by the subsampler 170. Pairs of fields, corresponding to the pairs selected by the time base changer 130, are selected by the subsampled time base changer and delay 180 from the subsampled video, to be used in the generation of motion vectors.

25 The time base changers 130 and 180 can operate according to synchronisation signals associated with the input video signal, the output video signal, or both. In the case in which only one synchronisation signal is supplied, the timing of fields of the other of the two video signals is generated deterministically within the time base changers 130, 180.

30 The pairs of fields of the subsampled input video signal selected by the

subsampled time base changer and delay 180 are supplied to a motion processor 185 comprising a direct block matcher 190, a correlation surface processor 200, a motion vector estimator 210, a motion vector reducer 220, a motion vector selector 230 and a motion vector post-processor 240. The pairs of input fields are supplied first to the direct block matcher 190 which calculates correlation surfaces representing the spatial correlation between search blocks in the temporally earlier of the two selected input fields and (larger) search areas in the temporally later of the two input fields.

From the correlation surfaces output by the block matcher 190, the correlation surface processor 200 generates a larger number of interpolated correlation surfaces, which are then passed to the motion vector estimator 210. The motion vector estimator 210 detects points of greatest correlation in the interpolated correlation surfaces. (The original correlation surfaces actually represent the difference between blocks of the two input fields; this means that the points of maximum correlation are in fact minima on the correlation surfaces, and are referred to as "minima"). In order to detect a minimum, additional points on the correlation surfaces are interpolated, providing a degree of compensation for the loss of resolution caused by the use of subsampled video to generate the surfaces. From the detected minimum on each correlation surface, the motion vector estimator 210 generates a motion vector which is supplied to the motion vector reducer 220.

The motion vector estimator 210 also performs a confidence test on each generated motion vector to establish whether that motion vector is significant above the average data level, and associates a confidence flag with each motion vector indicative of the result of the confidence test. The confidence test, known as the "threshold" test, is described (along with certain other features of the apparatus of Figure 1) in GB-A-2 231 749 and British Patent Application number 9307442.5.

A test is also performed by the motion vector estimator 210 to detect whether each vector is aliased. In this test, the correlation surface (apart from an exclusion zone around the detected minimum) is examined to detect the next lowest minimum. If this second minimum does not lie at the edge of the exclusion zone, the motion vector derived from the original minimum is flagged as being potentially aliased.

The motion vector reducer 220 operates to reduce the choice of possible motion vectors for each pixel of the output field, before the motion vectors are

supplied to the motion vector selector 230. The output field is notionally divided into blocks of pixels, each block having a corresponding position in the output field to that of a search block in the earlier of the selected input fields. The motion vector reducer compiles a group of four motion vectors to be associated with each block of the output field, with each pixel in that block eventually being interpolated using a selected one of that group of four motion vectors.

Vectors which have been flagged as "aliased" are re-qualified during vector reduction if they are identical, or similar within a selectable threshold, to non-flagged vectors in nearby blocks.

As part of its function, the motion vector reducer 220 counts the frequencies of occurrence of "good" motion vectors (i.e. motion vectors which pass the confidence test and the alias test, or which were re-qualified as non-aliased), with no account taken of the position of the blocks of the input fields used to obtain those motion vectors. The good motion vectors are then ranked in order of decreasing frequency. The most common of the good motion vectors which are significantly different to one another are then classed as "global" motion vectors. Three motion vectors which pass the confidence test are then selected for each block of output pixels and are supplied, with the zero motion vector, to the motion vector selector 230 for further processing. These three selected motion vectors are selected in a predetermined order of preference from:

(i) the motion vector generated from the corresponding search block (the "local" motion vector);

(ii) those generated from surrounding search blocks ("neighbouring" motion vectors); and

(iii) the global motion vectors.

The motion vector selector 230 also receives as inputs the two input fields which were selected by the subsampled time base changer and delay 180 and which were used to calculate the motion vectors, along with the control signal t . These fields are suitably delayed so that they are supplied to the motion vector selector 230 at the same time as the vectors derived from them. The motion vector selector 230 supplies an output comprising one motion vector per pixel of the output field. This motion vector is selected from the four motion vectors for that block supplied by the

motion vector reducer 220.

5 The vector selection process involves detecting the degree of correlation between test blocks of the two input fields pointed to by a motion vector under test. The motion vector having the greatest degree of correlation between the test blocks is selected for use in interpolation of the output pixel. A "motion flag" is also generated by the vector selector. This flag is set to "static" (no motion) if the degree of correlation between blocks pointed to by the zero motion vector is greater than a preset threshold.

10 The vector post-processor reformats the motion vectors selected by the motion vector selector 230 to reflect any vertical or horizontal scaling of the picture, and supplies the reformatted vectors to the interpolator 140. Using the motion vectors, the interpolator 140 interpolates an output field from the corresponding two (non-subsampled) interlaced input fields selected by the time base changer and delay 130, taking into account any image motion indicated by the motion vectors currently
15 supplied to the interpolator 140.

If the motion flag indicates that the current output pixel lies in a moving part of the image, pixels from the two selected fields supplied to the interpolator are combined in relative proportions depending on the temporal position of the output field with respect to the two input fields (as indicated by the control signal t), so that
20 a larger proportion of the nearer input field is used. If the motion flag is set to "static" then the temporal weighting is fixed at 50% of each input field. The output of the interpolator 140 is passed to an output buffer 150 for output as a high definition output signal, and to a down-converter 160 which generates a conventional definition output signal 165, using the motion flag.

25 The down-converter 160 allows a representation of the output of the apparatus (which may be, for example, a high definition video signal) to be monitored, transmitted and/or recorded using conventional definition apparatus. This has benefits because conventional definition recording equipment is significantly cheaper and very much more widespread than high definition equipment. For example, a simultaneous
30 output of conventional and high definition video may be required for respective transmission by terrestrial and satellite channels. Furthermore, if the output video signal is recorded directly onto film using, for example, an electron beam recorder,

the down-converter allows a simultaneous recording to be made onto video tape.

The subsampler 170 performs horizontal and vertical spatial subsampling of the input video fields received from the matrix 120, before those input fields are supplied to the time base changer 180. Horizontal subsampling is a straightforward operation in that the input fields are first prefiltered by a half-bandwidth low pass filter (in the present case of 2:1 horizontal decimation) and alternate video samples along each video line are then discarded, thereby reducing by one half the number of samples along each video line.

Vertical subsampling of the input fields is complicated by the fact that the input video signal 50 is interlaced. This means that successive lines of video samples in each interlaced field are effectively two video lines apart, and that the lines in each field are vertically displaced from those in the preceding or following field by one video line of the complete frame.

The construction and operation of the motion vector reducer 220 will now be described in more detail.

A block diagram of the motion vector reducer 220 is shown in Figure 2. The motion vectors, and their associated flags, derived by the motion vector estimator 210 are supplied via an input 300 of the motion vector reducer 220 to a ridge minima requalifier 310. The ridge minima requalifier 310 "repairs" motion vectors derived from elongate ("ridge") minima, by detecting whether another motion vector along the ridge direction has a matching coordinate at 90° to the ridge direction; if so, then that coordinate is assigned to the ridge vector and a "ridge flag" associated with that vector is reset. The ridge minima requalifier then supplies the vectors, and associated flags, to an alias vector validator 320 and a global vector detector 330. The alias vector validator 320 performs a test to check whether vectors indicated as potentially aliased can be justifiably requalified as non-aliased, and the vectors, and resulting associated flags, form one input to a global vector restricter 340. The global vector detector 330 derives a set of global motion vectors from the vectors supplied thereto. The global vectors output by the global vector detector 330 form a second input to the global vector restricter 340.

The global vector restricter 340 generates, for each global vector supplied thereto, a global mask array indicating an association between the global vector and

various blocks of the input image to which the local vectors supplied from the alias vector validator 320 correspond. The global mask arrays generated by the global vector restricter 340 are supplied to a block vector allocator 350 which also receives the (suitably delayed) local vectors, and associated flags, from the alias vector validator 320 and the global vectors from the global vector detector 330. The block vector allocator 350 assigns to each block of the input image to which the local vectors correspond a plurality of unique vectors which are then supplied via an output 360 of the motion vector reducer 220 to the motion vector selector 230.

The ridge minima requalifier 310, alias vector validator 320, global vector detector 330, global vector restricter 340 and block vector allocator 350 operate under the control (indicated schematically by dotted lines in the figure) of a microprocessor 370. The microprocessor 370 is also connected to the global vector detector 330 to provide for optional manual intervention in the global vector detection process to adjust various parameters. The adjusted parameters can be entered by an operator and supplied by the microprocessor 370 to the global vector detector 330. A two-way communication link 380 is also provided to connect the microprocessor 370 to a further, similar microprocessor (not shown). This is because the entire motion processor (indicated generally as 185 in Figure 1) may in fact be duplicated, the two motion processors performing the processing for alternate output fields. Thus, the vector reducer shown in Figure 2 may be duplicated, and the link 380 is provided for communication of the two microprocessors controlling the vector reduction operation.

The construction and operation of certain components of the motion vector reducer 220 will now be described in more detail.

The vectors, and flags, output by the ridge minima requalifier 310 are supplied to the global vector detector 330. The purpose of the global vector detector is to derive from the vectors supplied thereto a set of global motion vectors comprising a plurality of unique motion vectors selected from the most common of the local motion vectors. In general terms, the global vector detector 330 ranks local motion vectors corresponding to a given input image in order of decreasing frequency, and selects the eight (in this example) most common vectors as global vectors. However, it can be advantageous to apply various constraints to the process for selecting global vectors from the set of local vectors. In particular, even though large magnitude

(long) motion vectors (ie vectors representing rapid image motion) can be generated, it is desirable to prevent such long vectors from contributing to the global vectors. This is because the lack of spatial detail in fast-moving objects (eg due to camera integration blur) means that the operation of the motion vector selector 230 (Figure 1) for rapidly moving objects can be unreliable. It has been observed that rapidly moving portions of a video image have lower spatial detail than stationary or slowly moving portions. This means that even if a motion vector can be estimated which accurately represents the rapid motion of the object, the correlation test performed in the vector selector 230, which relies on detecting correlation between parts of the same object in two successive input fields, can be unreliable. It is therefore preferable to prevent such long vectors from being selected as global vectors to ensure that a large magnitude global vector cannot be incorrectly selected for interpolation of image areas where there is in fact only very small motion. This could produce subjectively disturbing effects in the output image.

A further constraint which is applied to the global vector detection process is to prevent a vector being accepted as a global vector unless it occurs at least a predetermined number of times over the input image as a whole. If a local motion vector is not sufficiently common, and thus represents motion of only a very small part of the image, it is preferable to eliminate the possibility of this vector being assigned to other parts of the image during the global vector detection stage.

A schematic block diagram of a global vector detector 330 in which the above constraints are implemented is shown in Figure 3. The input vectors, and associated flags, are supplied to an amplitude detector 400 which compares the amplitude, ie magnitude, of each input vector with a predetermined magnitude threshold 410 to decide whether the vector qualifies as a "long" vector. The magnitude threshold 410 may comprise a threshold for both the horizontal and vertical components of a motion vector. For example, the threshold may be set at half the maximum possible range, horizontally and vertically, of a motion vector. In this example, the maximum horizontal range of a motion vector is ± 64 pixels, and the maximum vertical range ± 32 pixels. Thus, the magnitude threshold 410 may be set at ± 32 pixels horizontally and ± 16 pixels vertically. If the magnitude of an input vector exceeds the magnitude threshold 410, the amplitude detector 400 sets a "long vector" flag

(LV).

As an alternative, the threshold 410 could represent the maximum value of the vector magnitude $\sqrt{x^2 + y^2}$, or a weighted sum to reflect a different scaling of the horizontal and vertical coordinates (e.g. $\sqrt{4x^2 + y^2}$ or $\sqrt{x^2 + 4y^2}$).

5 The input vectors supplied to the amplitude detector 400 are also supplied to a frequency array 420. Each entry in the frequency array 420 is uniquely addressed by a possible value of the local motion vectors. In this way, a local motion vector generated from the current pair of input fields is used to address the array and to increment the array entry corresponding to the value of the motion vector. However, 10 the long vector flags LV derived by the amplitude detector 400 are also supplied to the frequency array 420. If the long vector flag LV for a motion vector supplied to the frequency array 420 is set, then this prevents the associated motion vector from incrementing the frequency array.

15 After all the vectors (apart from long vectors) have been counted in this way by the frequency array 420, an array scanner 430 scans the frequency array 420 and identifies the eight array entries which show the highest count, for which the motion vectors differ from each other by at least a predetermined difference threshold 435. The addresses of these array entries represent the values of the eight most common local vectors (apart from long vectors) for the corresponding pair of input images. 20 The value of the count held by the frequency array for each of these eight vectors (that is, the frequency of occurrence of the vectors) is supplied by the array scanner 430 to a comparator 440 at the same time as the corresponding vector is supplied to an input of a switch 450. The comparator 440 compares the count, or frequency of occurrence, corresponding to each vector with a predetermined threshold number or 25 "global threshold", and provided the count exceeds the global threshold, the corresponding vector will be supplied to the output 460 of the switch 450 as a global motion vector. (A suitable value for the global threshold may vary depending on the particular image material being processed. By way of example, however, the global threshold may be set at 0.5% of the total number of local motion vectors derived 30 from an input image.) If the count corresponding to any of the eight most common vectors does not exceed the global threshold, indicating that the vector does not occur sufficiently frequently to qualify as a global vector, then the comparator 440 generates

an output to control the switch 450 such that the corresponding vector is not passed to its output 460 as a global vector, but is replaced by the zero motion vector. Assuming, however, that all eight of the most common vectors pass the global threshold test, then a set of eight unique global motion vectors will be supplied to the output 460 of the switch 450.

It will be appreciated that various further constraints could be applied in the global vector detector 330. In particular, any input local motion vectors for which the corresponding ridge or alias flags are set, or for which the confidence flag indicates that the vector failed the confidence test, may be prevented from contributing to the global vectors in a similar manner to long vectors, ie by preventing such vectors from incrementing the frequency array 420.

As indicated in Figure 3, the long vector flag LV derived for each of the input local motion vectors is supplied by the global vector detector 330, after an appropriate delay, to the block vector allocator 350 (Figure 2) for use in the subsequent vector reduction processing of the local vectors.

As shown in Figure 2, the global vector detector 330 is connected for communication with the microprocessor 370. This allows for the various thresholds used during global vector detection to be modified, if required, through manual intervention by an operator. The operator can enter new parameters, for example via a system controller (not shown), which are then supplied via the microprocessor 370 to the global vector detector 330.

The global vectors output by the global vector detector 330 are supplied to one input of the global vector restricter 340. At its other input, the global vector restricter 340 receives the processed local motion vectors and associated flags from the alias vector validator 320. The purpose of the global vector restricter 340 is to associate each of the global vectors with blocks of the input image at or near areas where motion identical or similar to that represented by the global motion vector was originally derived. This association will then be used by the block vector allocator 350 to prevent inappropriate global vectors being assigned to blocks, that is, to prevent a block from one area of the input image being assigned a global vector representing significantly different motion of a block in a remote part of the image.

The block vector allocator receives the LV flag from the amplitude detector

400, and can respond in a number of ways. In the embodiment described above, long vectors are already prevented from being used as global vectors; however, the block vector allocator 350 could in fact use the flag to prevent the use of long vectors at all. Such an arrangement is illustrated schematically in Figure 4, in which the LV flag associated with each motion vector controls a multiplexer 600 to pass either the associated vector or the zero vector to the input of the block vector allocator 350.

In another arrangement, illustrated schematically in Figure 5, the vector selection process for a particular block is restricted to the possible use of the long local vector for that block or the zero motion vector only. As shown in Figure 5, the LV flag for each block controls a multiplexer 610 to pass (to vector selection) either the three motion vectors assigned to that block during vector reduction or the local motion vector plus two zero motion vectors.

In another embodiment illustrated in Figure 6, a further amplitude detector 620, similar to the amplitude detector 400, is employed to compare each vector with a second, higher threshold value 630. Although the use of long vectors may in general be allowed (although restricted as described above), the use of vectors which are greater in size than this second threshold is not allowed. This is achieved by the use of a multiplexer 640 which passes each vector to the block vector allocator 350 only if that vector does not exceed the second threshold size 630.

In a further embodiment, the LV flag can be used to control the operation of the interpolator 140. This arrangement is illustrated in Figure 7, in which the LV flag controls a multiplexer 650, causing it to pass either the control signal t or a fixed value of 0.5 to the interpolator 140. This means that the use of temporal weighting (described above) can be inhibited when the LV flag is set. This can reduce disturbing image artifacts such as the visible judder caused by beating effects that may be observed on moving portions of "shuttered" images.

CLAIMS

1. Motion compensated video signal processing apparatus in which motion vectors are generated to represent image motion between a pair of input images of an input video signal, the apparatus comprising:

means for generating a plurality of local motion vectors to represent image motion of respective blocks of one input image of the pair between that image and the other image of the pair;

means for deriving from the local motion vectors a set of global motion vectors comprising a plurality of distinct motion vectors selected from the most common of the local motion vectors, each global vector having at least a predetermined number of occurrences among the local motion vectors; and

means for assigning to each of the blocks a group of motion vectors selected from the zero motion vector, the local motion vector for that block and the global motion vectors.

2. Apparatus according to claim 1, in which the means for deriving comprises means for preventing the selection, as global motion vectors, of two or more local motion vectors which differ from one another by less than a predetermined amount.

3. Apparatus according to claim 1 or claim 2, in which the predetermined number of occurrences can be adjusted under the control of an operator.

4. Apparatus according to any one of claims 1 to 3, comprising:

means for storing a frequency array comprising a plurality of addressable memory locations, each possible value of the motion vectors representing the address of a respective memory location;

means for incrementing a frequency value stored at the memory location corresponding to each of the plurality of motion vectors; and

means for detecting those memory locations having the highest stored frequency values.

5. Television standards conversion apparatus comprising apparatus according to any one of the preceding claims.
6. A method of motion compensated video signal processing in which motion
5 vectors are generated to represent image motion between a pair of input images of an input video signal, the method comprising the steps of:
- generating a plurality of local motion vectors to represent image motion of
 respective blocks of one input image of the pair between that image and the other
 image of the pair;
- 10 deriving from the local motion vectors a set of global motion vectors
 comprising a plurality of distinct motion vectors selected from the most common of
 the local motion vectors, each global vector having at least a predetermined number
 of occurrences among the local motion vectors; and
- 15 assigning to each of the blocks a group of motion vectors selected from the
 zero motion vector, the local motion vector for that block and the global motion
 vectors.

7. Motion compensated video signal processing apparatus substantially as hereinbefore described with reference to the accompanying drawings.

5 8. A method of motion compensated video signal processing, the method being substantially as hereinbefore described with reference to the accompanying drawings.

9. Television standards conversion apparatus substantially as hereinbefore described with reference to the accompanying drawings.



Application No: GB 9706681.5
Claims searched: All

Examiner: Sue Willcox
Date of search: 23 April 1997

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.O): H4F FGM FEP FER FEX

Int Cl (Ed.6): G06T 7/20; H04N 5/14 7/01

Other: Online database: WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	GB 2231749 A SONY CORPORATION (see page 19, line 21 to page 21, line 9)	1,2,4 - 6

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.